Available online at www.sciencedirect.com

ScienceDirect

Procedia Engineering 159 (2016) 89 – 93

Procedia Engineering

www.elsevier.com/locate/procedia

Humanitarian Technology: Science, Systems and Global Impact 2016, HumTech2016, 7-9 June 2016, Massachusetts, USA

CrowdAR: a live video annotation tool for rapid mapping

Elliot Salisbury*, Sebastian Stein, Sarvapali Ramchurn

Agents Interaction and Complexity Research Group, Department of Electronics and Computer Science, University of Southampton

Abstract

Digital Humanitarians are a powerful and effective resource to analyse the vast amounts of data that disasters generate. Aerial vehicles are increasingly being used for gathering high resolution imagery of affected areas, but require a lot of effort to effectively analyse, typically taking days to complete. We introduce CrowdAR, a real-time crowdsourcing platform that tags live footage from aerial vehicles flown during disasters. CrowdAR enables the analysis of footage within minutes, can rapidly plot snippets of the video onto a map, and can reduce the cognitive load of pilots by augmenting their live video feed with crowd annotations.

Peer-review under responsibility of the Organizing Committee of HumTech2016

Keywords: real-time; crowdsourcing; search and rescue; UAV

1. Introduction

Major man-made and natural disasters have a significant and long-lasting economic and social impact on countries around the world. The response effort in the first few hours of the aftermath of the disaster is crucial to saving lives and minimising damage to infrastructure. In these conditions, emergency responders on the ground face a major challenge in trying to understand what is happening, where the casualties are, and how to get to them safely. In order to help support this response effort, satellite imagery and aerial footage are increasingly being used by emergency responders and charitable organisations (e.g., Rescue Global, Red Cross, etc..), to help rapidly map disasters. However, the vast amount of imagery coming from such sources takes significant time to analyse in order to determine rescue targets.

Given such challenges, in recent years, a new community of volunteers, called digital humanitarians [1] has emerged as a powerful and effective resource to help analyse and organise vast amounts of data coming from disasters. For example, Tomnod, the Standby Task Force, and MicroMappers are but a few organisations that exist to rally volunteers and to carry out rapid data analysis, using typical web-based tools, in the aftermath of major disasters.

To date, these efforts have been focused on annotating satellite and aerial imagery, tracing maps, or classifying tweets [1–4]. In particular, these tasks tended to be 'offline' tasks that could take days to complete. However, analyzing aerial imagery (e.g., from planes or UAVs) 'offline' has a number of drawbacks. First, if the aerial vehicle is flown manually, then the pilot must maintain continuous focused attention so that they can spot search targets (e.g., damaged buildings, floating wreckage in the ocean after the missing flight MH370) and then fly the vehicle to get a closer look, and to record footage of it for later crowdsourced analysis. Inevitably, the pilot will eventually become fatigued, as these tasks are typically tedious (i.e., flying over an ocean) or visually taxing (i.e., flying over a city), and thus crucial information may be missed [5]. Second, if the aircraft is flown on a fixed flight path, as is typical with UAVs, then the benefit of diverting off course when a target is identified (post-flight) is lost, which can enable responders to gather imagery of potentially important regions that lie at the boundary of the vehicles field of view. Third, in the past, typical crowdsourcing tools have been used to review aerial footage post-flight [6]. Responders who launched the UAV cannot get quantifiable results immediately. Thus, despite being located nearby, cannot

* Corresponding author.
E-mail address: e.salisbury@ecs.soton.ac.uk

1877-7058
Peer-review under responsibility of the Organizing Committee of HumTech2016
doi:10.1016/j.proeng.2016.08.069

know where best to focus their efforts straightaway. Using aerial footage for situational awareness can still be a long and costly process.

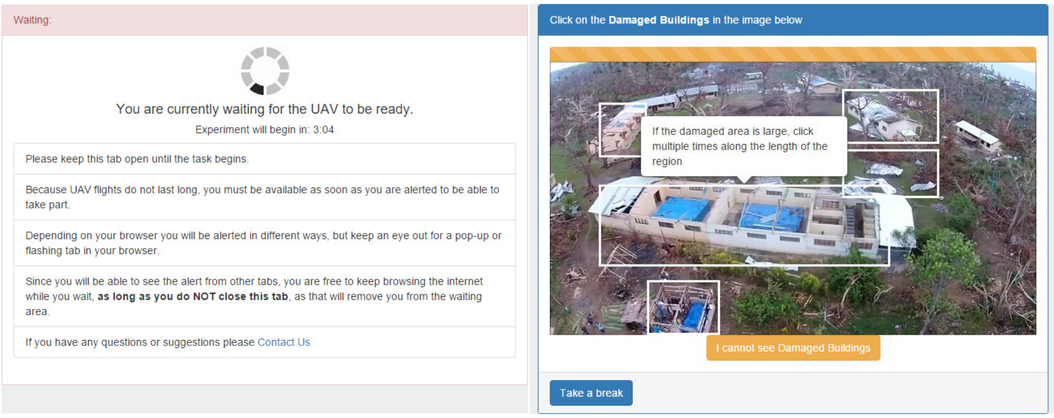

Fig. 1. The interfaces before the live broadcast begins. (a) Left: The Waiting Room; (b) Right: The Tutorial Page.

Automating these tasks would require artificial intelligence and computer vision algorithms. These algorithms need to be capable of reliably detecting any number of search targets, despite ever changing conditions (e.g., lighting conditions, orientation, partial occlusion). Work has been done to train computer vision classifiers for this purpose. For example, algorithms have been developed to detect damaged buildings [7], or injured people [8]. However, disasters are dynamic and often have different search requirements, it is infeasible to have an algorithm for each eventuality. Instead, a more general purpose approach would describe the target in natural language and use human recognition to understand and detect them.

Against this background, we propose Crowd Augmented Reality, or CrowdAR[1], a real-time crowdsourcing platform for annotating live video feeds, including, but not limited to, those from aerial vehicles [9]. CrowdAR uses scalable server-side processing for computer vision and real-time aggregation of crowd sourced data, and a specifically designed UI optimised for real-time crowdsourcing of live video. In the future, we intend to open-source CrowdAR, enabling everyone to use it.

Real-time crowdsourcing is the process of outsourcing work to multiple, simultaneously connected online workers [10]. When applied to disaster response, we can find and tag search targets in a live video feed. These annotations can then be aggregated and used to augment the live feed with the crowdsourced data, thereby reducing the cognitive load of the pilot. For example, informing a fatigued search and rescue pilot to floating wreckage or damaged buildings identified by the crowd, enabling the pilot to make more effective in-flight decisions. Furthermore, real-time crowdsourcing enables many more people to observe the live camera feed, and in so doing, they are more likely to observe and accurately identify important details than a single expert could. Additionally, while a single observer may become fatigued, a real-time crowdsourced team is a tireless workforce. More specifically, individuals of the crowd that tire and stop contributing to the task, can be replaced by new crowd members. As such, the pool of workers is always being replenished with fresh individuals.

Furthermore, through the use of in-flight analysis, CrowdAR reduces the delay between data gathering and the reviewing of footage. Moreover, in-flight analysis allows us to potentially react to events or features the crowd has spotted (e.g., damaged buildings, floating debris) and move the vehicle off its original search path to investigate.

To date, CrowdAR has been evaluated with both paid workers on crowdsourcing platforms, such as Amazon Mechanical Turk (AMT), and with digital humanitarian volunteers, such as MicroMappers, using videofeeds sourced from UAV flights after the Cyclone Pam disaster in Vanuatu, 2015.

2. CrowdAR

CrowdAR is a web-based tool used to crowdsource the annotation of a live video feed for supporting digital humanitarian efforts. CrowdAR can be used, for example, to rapidly map damaged buildings during a UAV flyby, or to identify search and rescue targets in both visually challenging or tedious environments.

CrowdAR consists of four main components, three from the perspective of the volunteers, and one final component that processes the volunteers data. The first stage that the digital humanitarians will encounter is the **Waiting Room**, here users wait for the live video feed to begin broadcasting. The second, and optional, stage is the **Tutorial** page, volunteers can acquaint

[1] https://crowdrobotics.org/projects/crowdar